# Open Compute Stack (OpenCS): a framework for parallelisation of equation-based simulation programs

#### 4 Dragan D. Nikolić

- 5 DAE Tools Project, Belgrade, Serbia, http://www.daetools.com
- 6 Corresponding author:
- 7 Dragan D. Nikolić
- 8 Email address: dnikolic@daetools.com

# ABSTRACT

In this work, the main ideas and the key concepts of the Open Compute Stack (OpenCS) 10 framework are presented. The framework provides a common platform for equation-based 11 modelling and parallel simulation on shared and distributed memory systems and allows 12 the same model specification to be used on different high performance computing systems 13 and architectures including heterogeneous systems. The main OpenCS components are 14 described: (1) platform-independent model specification data structures for description of 15 general systems of differential and algebraic equations, (2) a platform-independent method to 16 describe, store in computer memory and evaluate general systems of equations on diverse types 17 of computing devices, (3) an Application Programming Interface (API) for model specification, 18 parallel evaluation of model equations, simulation and model exchange, (4) algorithms for 19 partitioning of general systems of equations and inter-process data exchange, and (5) simulation 20 software for parallel numerical solution of general systems of differential and algebraic equations 21 on shared and distributed memory systems. The methodology and an API for the typical use 22 cases are presented. The benefits provided by the common modelling platform are discussed in 23 details such as: the model specification data structures providing a simple platform-independent 24 binary interface for model exchange, and the model equations stored as an array of binary 25 data which can be evaluated on virtually all computing devices with no additional processing. 26 The capabilities of the framework are illustrated using two large scale problems and the overall 27 performance and performance of individual phases of the numerical solution analysed. 28

# 29 INTRODUCTION

Equation-based mathematical modelling is one of the efficient methods for simulation of engi-30 neering problems described by a system of ordinary differential (ODE) or differential-algebraic 31 equations (DAE). On shared memory systems, the procedure for numerical solution of equation-32 based models includes the following computationally intensive tasks: (1) numerical integration 33 of the overall ODE/DAE system in time by a suitable solver (requires evaluation of model 34 equations), (2) linear algebra operations (mostly BLAS L1 vector operations and some of BLAS 35 L2 matrix-vector operations), (3) solution of systems of linear equations (requires evaluation 36 of derivatives), and (4) if requested, integration of sensitivity equations (requires evaluation of 37 sensitivity residuals). On distributed memory systems, every processing element (PE) performs 38

the same tasks on one part of the overall system (ODE/DAE sub-sybstem) and an inter-process
data exchange required by the linear algebra and equation evaluation functions. In general,
simulation programs for this class of problems are developed using:

- General-purpose programming languages such as C/C++ or Fortran and one of available
   suites for scientific applications such as SUNDIALS (Hindmarsh et al., 2005), Trilinos
   (Heroux et al., 2005) and PETSC (Balay et al., 2015)
- 45 2. Modelling languages such as Ascend (Piela et al., 1991), gPROMS (Barton and Pantelides,
   46 1994), APMonitor (Hedengren et al., 2014) and Modelica (Fritzson and Engelson, 1998)
- Multi-paradigm numerical languages such as Matlab (The MathWorks, Inc., 2018a), Math ematica (Wolfram Research, Inc., 2015) and Maple (Waterloo Maple, Inc., 2015)
- 4. Higher-level fourth-generation languages (Python) and modelling software such as DAE
   Tools (Nikolić, 2016) and Assimulo (Andersson et al., 2015)
- 5. Libraries for finite element (FE) analysis and computational fluid dynamics (CFD) such
   as deal.II (Bangerth et al., 2007), libMesh (Kirk et al., 2006), and OpenFOAM (The
   OpenFOAM Foundation, 2018)
- Computer Aided Engineering (CAE) software for finite element analysis and computational
   fluid dynamics such as HyperWorks (Altair, 2018), STAR-CCM+ and STAR-CD (Siemens,
   2018), COMSOL Multiphysics (COMSOL, Inc., 2018), ANSYS Fluent/CFX (Ansys, Inc.,
   2018) and Abaqus (Dassault Systemes, 2018)

A detailed discussion of capabilities and limitations of the available approaches for speci-58 fication of model equations and development of large-scale simulation programs are given in 59 Nikolić (2016, 2018, 2019). In all approaches, an interface to a particular ODE/DAE solver 60 must be implemented to provide the information required for numerical integration in time (Fig. 61 1). The solver interface is directly implemented in general-purpose programming languages 62 (i.e. as user-supplied functions). In other approaches, the solver interface is built around the 63 internal simulator-specific data structures representing the model. For instance, the source code of 64 modelling languages is typically parsed into an Abstract Syntax Tree (AST). The produced AST 65 can be transformed into a simulator-specific data structure or used to generate C source code as 66 in OpenModelica (Fritzson et al., 2005) and JModelica (Akesson et al., 2010). Other modelling 67 software such as DAE Tools use the operator overloading technique to produce a tree-like data 68 structure (Evaluation Tree). CAE software perform a discretisation of Partial Differential Equa-69 tions (PDE) on a specified grid: (a) on unstructured grids, the results of discretisation using the 70 Finite Element (FE) or Finite Volume (FV) methods are the mass and stiffness matrices and load 71 vectors, and (b) on structured grids, the results of discretisation using the Finite Difference (FD) 72 method are the stencil data (nodes arrangement and their coefficients). The simulator-specific data 73 structures, sparse matrix-vector (SpMV) and matrix-matrix (SpMM) operations or stencil codes 74 are then utilised by the ODE/DAE solver interface to evaluate model equations and derivatives. 75 The idea in this work is to separate a high-level (simulator-dependent) model specification 76 procedure, typically performed only once, from its parallel (in general, simulator-independent) 77 numerical solution. While description of models and generation of a system of equations can be 78 performed in many different ways depending on the type of the problem and the method applied 79 by a simulator, the numerical solution procedure always requires the same (low-level) information. 80 For instance, a high-level model specification for the problems governed by partial differential 81 equations can be created using a modelling language or a CAE software. The low-level model 82 description is internally generated by simulators utilising various discretisation methods and 83 results in a system of differential equations (ODE or DAE). However, the information required for 84



**Figure 1.** An overview of available modelling approaches. A high-level model description is created using the modelling or general purpose programming languages, FEA/CFD libraries and CAE software. A low-level model description is generated in a problem and simulator specific way and used to implement an interface to ODE/DAE solvers. The solver interface utilises the simulator-specific data structures to provide the information required for integration of the ODE/DAE system in time.

numerical solution in both cases are essentially identical: the data about the number of variables, 85 their names, types, absolute tolerances and initial conditions, and the functions for evaluation 86 of equations and derivatives. Therefore, the low-level model description coupled with a method 87 for parallel evaluation of model equations on different computing devices can be a basis for 88 a universal software for parallel simulation of general systems of differential equations on all 80 important platforms. In general, such a model description, due to its simplicity, can be generated 90 and utilised by any existing simulator. This way, simulations can be performed on platforms not 91 supported by that particular simulator or the simulation performance on the supported platforms 92 can be improved by evaluating model equations in parallel on devices that are not currently utilised. 93 In addition, the same platform-independent model description can be used for model exchange 94 and benchmarks between different simulators, solvers, individual computing devices and high 95 performance computing platforms (i.e. between heterogeneous clusters, where evaluation of 96 model equations is currently not available for different architectures). An efficient evaluation of 97 model equations is of utmost importance. For instance, very often more than 85% of the total 98 integration time is spent on evaluation of equations and derivatives (Nikolić, 2018). Since most of 90 the modern computers and many specially designed clusters are equipped with additional stream 100 processors/accelerators such as Graphics Processing Units (GPU), Field Programmable Gate 101 Arrays (FPGA) and manycore processors (Xeon Phi), the simulation software must be specially 102 designed to effectively take advantage of multiple architectures. While parallel evaluation of 103 model equations on general purpose processors is fairly straightforward and different techniques 104 are applied by different simulators, evaluation on streaming processors is rather difficult. Stream 105 computing differs from traditional computing in that the system processes a sequential stream of 106 elements: a kernel is executed on each element of the input stream and the result stored in an 107

output stream. Thus, the data structures representing the model equations must be designed to 108 support evaluation on both systems (often simultaneously in heterogeneous computing setups). 109 To this end, the Open Compute Stack (OpenCS) framework has been develop to provide: 110 1. Model specification data structures for a platform-independent description of general 111 **ODE/DAE** systems of equations 112 2. A platform-independent method to describe, store in computer memory and evaluate 113 general systems of equations of any size on diverse types of computing devices 114 3. An Application Programming Interface (API) for model specification, parallel evaluation 115 of model equations, model exchange and a generic interface to ODE/DAE solvers 116 4. Algorithms for partitioning of general systems of equations and inter-process data exchange 117 (for simulations on distributed memory systems) 118 5. Simulation software for parallel numerical solution of general ODE/DAE systems of 119 equations on shared and distributed memory systems 120 This way, the OpenCS framework offers a common platform for specification of equation-based 121 models, parallel evaluation of equations on diverse types of computing devices, model exchange 122 and parallel simulation of large-scale systems of differential equations on shared and distributed 123 memory systems. 124 OpenCS is free software released under the GNU Lesser General Public Licence. The 125 installation packages, compilation instructions and more information about the OpenCS software 126 can be found on the DAE Tools website (http://www.daetools.com/opencs.html). The source code 127 is available from the SourceForge subversion repository: https://sourceforge.net/p/daetools/code 128 and located in the trunk/OpenCS directory. 129 The framework is based on the methodology for parallel numerical solution of general systems 130 of non-linear differential and algebraic equations on heterogeneous and distributed memory 131 systems presented in Nikolić (2018, 2019). In the OpenCS approach, the model specification 132 contains only the low-level information directly required by solvers. The model equations are 133 transformed into the Reverse Polish (postfix) notation and stored as an array of binary data 134 (a Compute Stack) for direct evaluation on all platforms with no additional processing nor 135 compilation steps. The OpenCS model specification, represented by a Compute Stack Model, 136 provides a common interface to ODE/DAE solvers and can be generated using the OpenCS API 137 in two ways (Fig. 2): (a) direct implementation in C++ and (b) export of existing models from 138 third-party simulators. Individual equations (Compute Stacks) are evaluated by a stack machine 139 (Compute Stack Machine) using the Last In First Out (LIFO) queues. Systems of equations are 140 evaluated in parallel using a Compute Stack Evaluator interface which manages the Compute 141 Stack Machine kernels. Two APIs/frameworks are used for parallelism: (a) the Open Multi-142 Processing (OpenMP) API for parallelisation on general purpose processors (multi-core CPUs, 143 Xeon Phi), and (b) the Open Computing Language (OpenCL) framework for parallelisation on 144 streaming processors (GPU, FPGA) and heterogeneous systems (CPU+GPU, CPU+FPGA). 145



**Figure 2.** The OpenCS modelling approach. The (low-level) model specification is created using the OpenCS API and stored in a Compute Stack Model data structure which provides a generic interface to ODE/DAE solvers. Model equations, transformed into the postfix notation and stored as an array of binary data, are evaluated using the Compute Stack Machine kernels managed by a Compute Stack Evaluator.

A generic simulation software has been provided by the framework to utilise the low-level 146 information stored in Compute Stack Models (Nikolić, 2019). Simulations can be executed 147 sequentially on a single processor or in parallel on message passing multiprocessors, where every 148 processing element integrates one part (sub-system) of the overall ODE/DAE system in time and 149 performs an inter-process communication between the processing elements (Fig. 3). Simulation 150 inputs are specified in a generic fashion as files in a (platform independent) binary format. The 151 input files are generated using the OpenCS API (one set per processing element) and contain the 152 serialised model specification data structures and solver options. This way, the OpenCS model 153 specification stored in input files is used as a simple binary interface for model exchange. This 154 approach differs from the typical model-exchange/co-simulation interfaces in that it does not 155 require a human or a machine readable model definition as in modelling and model-exchange 156 languages such as Modelica and CellML (https://www.cellml.org) nor a binary interface (C API) 157 implemented in shared libraries as in Simulink (The MathWorks, Inc., 2018b) and Functional 158 Mock-up Interface (https://www.fmi-standard.org). For instance, the model equations in OpenCS 159 are specified as an array of binary data for direct evaluation on all platforms with no additional 160 processing steps. However, it must be kept in mind that the main purpose is an exchange of 161 individual large-scale models whose equations can be evaluated on different computing devices 162 and which can be simulated on different high-performance computing platforms. Although 163 technically possible, use of OpenCS models as building blocks in other simulators is not the 164 major goal of OpenCS. 165



Figure 3. Parallel simulation on distributed memory systems using the OpenCS framework

The OpenCS framework offers the numerous benefits. A single software is used for numerical 166 solution of any system of differential and algebraic equations (ODE or DAE) of any size and 167 on all platforms. The model specification contains only the low-level model description and 168 therefore can be generated from any modelling software. The model specification data structures 169 are stored as files in a binary format and used as inputs for parallel simulations on all platforms. 170 Model equations are specified in a platform and programming language independent fashion 171 as an array of binary data. Equations of any type (differential or algebraic) and any size are 172 supported and can be evaluated on virtually all computing devices (including heterogeneous 173 systems). Switching to a different computing platform for evaluation of model equations is 174 straightforward and controlled by an input parameter. 175

OpenCS model description is created using the OpenCS API in two ways: direct implementation in C++ or export of existing models from third-party simulators. The most important use-cases scenarios of the OpenCS framework include:

179 1. Universal parallel simulations on shared and distributed memory systems

- Parallel evaluation of model equations (i.e. in simulators with no support for parallel
   evaluation or using the computing devices which are currently not utilised)
- <sup>182</sup> 3. Model-exchange

4. Use as a simulation engine behind Modelling or Domain Specific Languages

In addition, since the common model-specification in a binary format is used on all platforms, OpenCS models can be used for benchmarks between different simulators, ODE/DAE solvers, individual computing devices (i.e. to compare memory and computation performance during evaluation of equations) and high performance computing systems. For example, benchmarks between heterogeneous CPU+GPU and CPU+FPGA clusters are now possible without re-implementation of the model for a completely different architecture: in the OpenCS approach, the same data are used to evaluate equations on all computing devices.

The article is organised in the following way. First, the methodology, key concepts, data structures, API and implementations are presented. Next, the typical use-case scenarios accompanied with the sample ODE/DAE problems are analysed. Finally, a summary of the most important capabilities of the OpenCS framework and directions for future work are given in the last section.

# 195 **METHODS**

<sup>196</sup> The framework is based on the methodology for parallel numerical solution of general systems of

<sup>197</sup> differential and algebraic equations on heterogeneous and distributed memory systems presented

<sup>198</sup> in Nikolić (2018, 2019). The methodology consists of the following parts:

- A method for transformation of model equations into a data structure suitable for parallel
   evaluation on different computing platforms (the Compute Stack approach)
- 201 2. Data structures for model specification
- <sup>202</sup> 3. An algorithm for partitioning of general systems of equations
- 4. An algorithm for inter-process data exchange
- 5. Simulation software for integration of general ODE/DAE systems in time

#### <sup>205</sup> The key concepts and data structures

- <sup>206</sup> The methodology is based on several concepts, each providing a distinct functionality:
- Compute Stack The Reverse Polish (postfix) notation expression stack used as a platform
   and programming language independent method to describe, store in computer memory
   and evaluate equations of any type and any size (Nikolić, 2018). Compute Stacks are
   automatically generated from equations in infix notation using the OpenCS API.
- Compute Stack Machine A stack machine used to evaluate a single equation (that is a single
   Compute Stack) using LIFO queues (function *evaluateComputeStack* in the supplemental
   source code listing S1).
- Compute Stack Evaluator An interface for parallel evaluation of systems of equations (*csComputeStackEvaluator\_t* class in the supplemental source code listing S2).
- Compute Stack Model Data structure that serves as the main storage for the model specification
   and includes the information required for numerical solution, either sequentially or in
   parallel (*csModel\_t* data structure in the supplemental source code listing S3). In general,
   it can be used to describe a system of equations of any type.
- Compute Stack Differential Equations Model An abstract class that provides: (a) a model
   exchange interface, and (b) a generic interface to ODE/DAE solvers (*csDifferentialEquationModel\_t* class in the supplemental source code listing S4).
- Compute Stack Simulator Software for sequential and parallel simulation of general ODE and
   DAE systems in time (*csSimulator*).
- Compute Stack Model Builder An interface that provides an API for model specification, an
   algorithm for partitioning of general systems of equations with multiple load balancing
   constraints, generation and export of Compute Stack models (*csModelBuilder\_t* class in
   the supplemental source code listing S5).

**Compute Stack Number** A user-defined Real number class for creation of mathematical ex-229 pressions representing the model equations (csNumber\_t class in the supplemental source 230 code listing S6). Equations are specified in infix notation and transformed into the Compute 231 Stacks using the operator overloading technique. It is based on the same principles as the 232 Evaluation Tree data structure described in Nikolić (2018), and provides the following 233 functionality: (a) standard mathematical operations and functions (re-defined to operate on 234 the csNumber\_t objects), (b) evaluation of equations, (c) export into the LaTex format, and 235 (d) generation of Compute Stack arrays. 236

Compute Stack Graph Partitioner An interface for partitioning of graphs utilised by the partitioning algorithm in the Model Builder (*csGraphPartitioner\_t* in the supplemental source code listing S7).

#### 240 Model equations

In the Compute Stack approach, model equations are transformed into the Reverse Polish (postfix) 24 notation (Nikolić, 2018). Each mathematical operation and its operands are described by a 242 specially designed *csComputeStackItem* t data structure and every equation is transformed into an 243 array of these structures (a Compute Stack). In general, any type of expressions involving standard 244 mathematical operators and functions, numerical constants, variables and their derivatives are 245 supported (linear or non-linear, algebraic or differential equations). Most of the functions from C 246 numerics library (the <math.h> header) are available such as: unary (+, -) and binary operators 247 (+, -, \* and /), and unary and binary functions (sqrt, pow, log, log10, exp, min, max, floor, ceil, 248 abs, sin, cos, tan, asin, acos, atan, sinh, cosh, tanh, asinh, acosh, atanh, atan2 and erf). 249 Individual equations are evaluated by a stack machine (Compute Stack Machine) using a Last 250 In First Out queue. Due to its simplicity, equations can be evaluated on virtually all computing 251

devices (Nikolić, 2018). Since the model equations are stored as an array of binary data they
can be directly evaluated on all platforms with no additional processing nor compilation steps.
An overview of the Compute Stack Machine and the required input data are given in Fig. 4. As

inputs, it requires the evaluation context object with the run-time information, a Compute Stack

- array as a stream of contiguous data and random access data arrays with variable values (x), time
- derivatives (dx/dt) and degrees of freedom (y).



**Figure 4.** The Compute Stack Machine for evaluation of general  $F(x, \frac{dx}{dt}, y)$  expressions specified as Compute Stacks. *x*, dx/dt and *y* are arrays with variable values, time derivatives and degrees of freedom, respectively. The inputs are the evaluation context object with the run-time information, a Compute Stack array as a stream of contiguous data and three random access arrays (*x*, dx/dt and *y*).

Systems of equations are stored in memory as a single one-dimensional array of *csComputeStackItem\_t* objects populated with Compute Stacks from all equations. Parallel evaluation of systems of equations is performed through a common interface called a Compute Stack Evaluator. Two implementations are available: (a) the OpenMP API is used for parallelisation on general purpose processors, and (b) the OpenCL framework is used for parallelisation on streaming processors and heterogeneous systems. In the OpenMP implementation, every thread evaluates a chunk of the total number of equations, one at the time. In the OpenCL implementation, every
 work-item evaluates only a single equation. Each thread/work-item executes a for loop where
 mathematical operations are performed in a single IF block controlled by the type of mathematical
 operation.

#### 268 Model specification data structures

In the OpenCS approach, the model specification contains only the low-level information directly 269 required by ODE/DAE solvers, stored in the *csModel t* data structure (Nikolić, 2019). For 270 sequential simulations, the system is described by a single *csModel t* object. For parallel 271 simulations, the system is described by an array of *csModel\_t* objects each holding information 272 about one ODE/DAE sub-system. Every model contains the following data: (a) the model 273 structure with the information about the variable names, types, absolute tolerances and initial 274 conditions: csModelStructure\_t structure, (b) the model equations represented by Compute Stack 275 arrays: csModelEquations t structure, (c) the sparsity pattern of the ODE/DAE (sub-) system 276 (required for evaluation of derivatives): csSparsityPattern\_t structure, (d) the partition data 277 (used for inter-process communication): csPartitionData t structure, and (e) the Compute Stack 278 evaluator instance: *csComputeStackEvaluator\_t* object. Models are created using the Compute 279 Stack Model Builder that provides an API for model specification and hides the implementation 280 details of the OpenCS framework. 281

Model exchange capabilities and an interface to ODE/DAE solvers are provided by the *csDifferentialEquationModel\_t* class. It contains an instance of the *csModel\_t* class and functions for loading of models from input files, retrieving the information and the sparsity pattern of the ODE/DAE system, setting the variable values/derivatives, exchanging the adjacent variables among the processing elements using the MPI interface, and evaluating equations and derivatives.

#### 287 Partitioning of general systems of equations

Large-scale numerical simulations on parallel computers require the distribution of equations 288 among the processing elements so that the duration of each phase of the numerical solution is 289 approximately the same. Therefore, the workload (storage and computation) in each phase and 290 the inter-process communication volume must be well balanced among the processing elements 291 for maximum performance. Computationally the most intensive phases of the numerical solution 292 are: (1) evaluation of equations, (2) solution of systems of linear equations, and (3) evaluation 293 of derivatives. Combined, they amount to more than 95% of the total integration time (Nikolić, 294 2018). Since it is critical that every processor have an equal amount of work from each phase of 295 the computation, the multiple quantities must be load balanced simultaneously. 296

The algorithm for partition of general systems of equations is described in Nikolić (2019). In 297 the OpenCS framework, this algorithm is improved and re-implemented in C++ (for performance 298 reasons). In the original algorithm, a graph of the ODE/DAE system is constructed and always 299 partitioned using the METIS library (Karypis and Kumar, 1995). In this work, the graph 300 partitioning is performed by the Compute Stack Graph Partitioner interface (csGraphPartitioner\_t 301 class) and separated from the main algorithm. This way, the algorithm can support the user-defined 302 graph partitioners to exploit a problem-specific structure of model equations. At the moment, 303 the following graph partitioner implementations are available: (1) Simple graph partitioner 304 (csGraphPartitioner\_Simple) - splits a graph into the specified number of partitions with no 305 load balancing analysis (i.e. used for generation of Compute Stack models for sequential 306 simulations), and (2) Metis graph partitioner (csGraphPartitioner\_Metis) - partitions the graphs 307

into a user-specified number k of parts using either the Multilevel k-way partitioning paradigm 308 or the Multilevel recursive bisectioning paradigm implemented in METIS. (3) 2D-Npde graph 309 partitioner (csGraphPartitioner\_2D\_Npde) - partitions the specified number of partial differential 310 equations (Npde) distributed on a uniform two-dimensional grid by dividing the grid into the 311 requested number of regions. The partitioning algorithm applies a static load balancing method. 312 The workloads can be accurately and precisely estimated by taking into consideration several 313 properties of equations and partitions. The partition properties used by the algorithm are: number 314 of equations  $(N_{eq})$ , number of adjacent variables  $(N_{adj})$ , number of items in the Compute Stack 315 array  $(N_{cs})$ , number of non-zero items in the partition's incidence matrix  $(N_{nz})$ , number of 316 floating point operations (FLOPs) required for evaluation of equations  $(N_{flops})$ , and number of 317 FLOPs required for evaluation of derivatives  $(N_{flops})$ . The memory and computation workloads 318 in individual phases can be estimated using the partition properties as discussed in (Nikolić, 319 2019).  $N_{cs}$ ,  $N_{nz}$ ,  $N_{flops}$  and  $N_{flops_i}$  can be specified as additional balancing constraints for the 320 graph partitioning. In particular, the number of FLOPs required for evaluation of equations 321 and derivatives, that is the computation load, can be very accurately estimated by analysing 322 the Compute Stack arrays. Moreover, the partitioning algorithm accepts a pair of dictionaries 323 specifying the number of FLOPs for individual unary and binary mathematical operations (Nikolić, 324 2019). For instance, evaluation time of trigonometric functions on a traditional CPU is different 325 from the evaluation time on a GPU. Thus, the algorithm can produce the load balanced partitions 326 for diverse types of computing devices. 327

Partitioning of systems of equations in some cases is problem-specific and the generic graph 328 partitioners often produce partitions with the excellent balance of workloads but poor overall 329 simulation performance. The reason for this is the structure of partitions resulting in inefficient 330 preconditioners and a high number of iterations to reach convergence in the linear solver, as 331 discussed in Nikolić (2019). Therefore, custom user-defined partitioners are required to take 332 advantage of a problem-specific structure of model equations (i.e. the systems produced by 333 discretisation of well known partial-differential equations on uniform grids and the systems which 334 require the coupled treatment of all differential equations to ensure conservation such as the 335 compressible Euler and incompressible Navier-Stokes equations). 336

#### 337 Inter-process data exchange

Numerical solution on distributed memory systems requires an inter-process communication routine for exchange of adjacent unknowns (unknowns that belong to other processing elements). The algorithm for data exchange among processing elements is simple and only the point-to-point communication routines are required (Nikolić, 2019). It is fully generic and utilises the data resulting from the partitioning algorithm stored in the *csPartitionData\_t* data structure.

#### **Generic simulation software**

The OpenCS framework provides a simulator for integration of general systems of differential equations in time (*csSimulator*) which can simulate both ODE and DAE systems (Nikolić, 2019). The simulator is cross-platform and can be executed sequentially on a single processor or in parallel on message passing multiprocessors. Simulation inputs are specified in a platformindependent way using input files with the model specification and run-time options. This way, the same model can be simulated using the same software on all platforms. An overview of the solution procedure on shared memory systems is given in Fig. 5. The

solution process consists of: (1) numerical integration in time, (2) linear algebra operations, (3)

solution of systems of linear equations (in general, iterative methods are used for large scale
systems), (4) (optionally) integration of sensitivity equations. The Compute Stack Evaluator is
utilised by the Compute Stack Model for parallel evaluation of equations residuals (DAE systems)
or the right hand side (ODE systems), and for evaluation of derivatives required for computation
of the preconditioner and integration of sensitivity equations. Depending on the simulation
options, the Compute Stack Evaluator can utilise a single or multiple computing devices.

The parallel solution on distributed memory systems requires the same tasks, but applied to integration of only one part of the overall system (ODE/DAE sub-system). Therefore, the software for numerical solution on shared memory systems is used as the main building block for distributed memory systems as depicted in Fig. 6. The additional functionality that is required includes: (a) an inter-process communication routine for exchange of adjacent unknowns, and (b) linear algebra routines for distributed memory systems (already available from the SUNDIALS suite). Both routines are implemented using the MPI C interface.

For integration of DAE systems in time the software uses the variable-step variable-order 365 backward differentiation formula available in SUNDIALS IDAS solver (Hindmarsh et al., 2005). 366 For integration of ODE systems in time the software uses the variable-step variable-order Adams-367 Moulton and backward differentiation formulas available in SUNDIALS CVodes solver (Serban 368 and Hindmarsh, 2005). Systems of linear equations are solved using the Krylov-subspace iterative 369 methods. At the moment, the generalised minimal residual solver from the SUNDIALS suite is 370 available. Both solvers utilise preconditioners available from the Trilinos suite (Heroux et al., 371 2005): IFPACK, ML and AztecOO built-in preconditioners. Evaluation of model equations and 372 derivatives is performed through the Compute Stack Evaluator interface. 373

Simulation inputs are specified using the data files with the serialised model specification 374 data structures. The list of input files (one set for every processing element) is given in Table 375 1. PE in file names is an integer identifying the processing element equal to the value returned 376 from MPI\_Comm\_rank function. For sequential simulations a single set of input files is required. 377 Each file contains a serialised data structure member of the *csModel t* class: *csModelStructure t*, 378 csModelEquations\_t, csSparsityPattern\_t and csPartitionData\_t. While the model specification 379 remains unaltered, simulations can be performed for different time horizons, different solver and 380 preconditioner options and using different computing devices for evaluation of model equations. 381 Thus, the simulation options are specified in a human readable JSON format. and contain 382 four sections: "Simulation" (run-time data), "Model" (ODE/DAE model options), "Solver" 383 (options for the ODE/DAE solver) and "LinearSolver" (the linear solver and the preconditioner 384 options). Names of the solver/preconditioner parameters are identical to the original names used 385 by the corresponding libraries or to the names of *Set* functions (i.e. the *MaxOrd* parameter 386 specified using the IDASetMaxOrd function in the SUNDIALS suite). The typical content of the 387 simulation\_option.json file for ODE and DAE problems are given in the supplemental source 388 code listings S8 and S9, respectively. 389

Simulation results are saved in Comma Separated Value (.csv) format into the output directory specified by the *Simulation.OutputDirectory* option. In addition, the detailed solvers statistics is generated for every processing element and saved in JSON format into the output directory.



Figure 5. OpenCS simulation on shared memory systems



Figure 6. OpenCS simulation on distributed memory systems

**Table 1.** Input data files for OpenCS simulations.

Input file	Contents
model_structure-[PE].csdata	Serialised <i>csModelStructure_t</i> data structure
model_equations-[PE].csdata	Serialised csModelEquations_t data structure
sparsity_pattern-[PE].csdata	Serialised csSparsityPattern_t data structure
partition_data-[PE].csdata	Serialised csPartitionData_t data structure
simulation_options.json	Simulation, DAE and linear solver parameters

# **APPLICATION PROGRAMMING INTERFACE**

The OpenCS framework provides an Application Programming Interface for model specifi-394 cation and typical use case scenarios such as (1) parallel evaluation of model equations, (2)395 model exchange, (3) simulation on shared memory systems, and (4) simulation on distributed 396 memory systems. The key concepts of the OpenCS framework and the corresponding API are 397 implemented in the following libraries: (1) cs machine.h (header-only Compute Stack Machine 398 implementation in C99), (2) libOpenCS\_Evaluators (sequential, OpenMP and OpenCL Compute 399 Stack Evaluator implementations), (3) libOpenCS Models (Compute Stack Model, Compute 400 Stack Differential Equations Model and Compute Stack Model Builder implementations), (4) 401 libOpenCS Simulators (Compute Stack ODE and DAE Simulator implementations) and a stan-402 dalone simulator *csSimulator* (for both ODE and DAE problems). The framework internally 403 utilise computing devices for evaluation of model equations and performs file system I/O oper-404 ations and inter-process communication using the MPI interface (in parallel simulations). The 405 structure and the main components of the framework are illustrated in Fig. 7. 406



Figure 7. The structure and the main components of the OpenCS framework

#### 407 Model specification

- <sup>408</sup> In the OpenCS framework, models are developed using the Model Builder interface (*csModel*-
- <sup>409</sup> *Builder\_t* class). The model equations can be specified in C++ application programs or exported
- 410 from existing models in third party simulators. The procedure is identical for all models in both
- cases. For simulations on shared memory systems it includes the following steps (source codelisting 1):
- Step 1. Initialise the model builder with the number of variables and the number of degrees of
   freedom. Other options such the default variable value, absolute tolerance, and variable
- name can be optionally set.

their time derivatives and degrees of freedom. *csNumber\_t* is a user-defined Real number 417 class providing all standard mathematical operators and functions as in the C numerics 418 library (<math.h> header). This way, the equations are specified in the same way as in 419 C/C++. 420 **Step 3.** Set the initial conditions. For DAE problems a consistent set of initial conditions will be 421 calculated before simulation. 422 **Step 4.** Generate Compute Stack models by partitioning the system of equations. For simulations 423 on shared memory systems the model is partitioned into a single partition. For simulations 424 on message passing multiprocessors the system is partitioned into a specified number of 425 partitions (one per processing element). Typically, the Simple graph partitioner is used by 426 the partitioning algorithm for sequential simulations. 427 **Step 5.** Use the generated model(s) directly or export them into a specified directory. 428 For parallel simulations on distributed memory systems the procedure is identical. The only 429 difference is in the Step 4. where the system is partitioned using the METIS or a user-defined 430 graph partitioner. The graph partitioning procedure for parallel simulations includes the following 431 steps (source code listing 2): 432 **Step 4.2** Instantiate METIS or a user-defined graph partitioner. In Metis, two algorithms are 433 available: (1) PartGraphKway (Multilevel k-way partitioning algorithm), and (2) Part-434 GraphRecursive (Multilevel recursive bisectioning algorithm). Optionally, change the 435 default options of the partitioning algorithm. 436 **Step 4.2.1** Specify the load balancing constraints. Four additional balancing constraints 437 are available:  $N_{cs}$ ,  $N_{nz}$ ,  $N_{flops}$  and  $N_{flops}$  *i*. For example,  $N_{cs}$  and  $N_{flops}$  can be used 438 to balance the memory load (proportional to the number of Compute Stack items, 439  $N_{cs}$ ) and the computation load for evaluation of model equations (proportional to the 440 number of FLOPs for evaluation of equations,  $N_{flops}$ ). 441 **Step 4.2.2** Set the graph partitioner options (METIS specific). 442 **Step 4.2.3** By default, the partitioning algorithm assumes that all mathematical operations 443 require a single FLOP. This behaviour can be changed by specifying a pair of 444 dictionaries with a number of FLOPs for individual mathematical operations: (1) 445 unaryOperationsFlops for unary operators (+, -) and functions (sqrt, log, log10, exp, 446 sin, cos, tan, ...), and (2) binaryOperationsFlops for binary operators (+, -, \*, /) and 447 functions (pow, min, max, atan2). If a mathematical operation is not in the dictionary, 448 it is assumed that it requires a single FLOP. This way, the total number of FLOPs 449

**Step 2.** Create model equations using the provided *csNumber* t objects representing variables,

416

- 450 can be accurately estimated for every computing device.
- 451 Step 4.3 Partition the system into the specified number of processing elements (Npe).

**Listing 1.** Model specification procedure for simulation on shared memory systems (DAE problem)

```
452
453
    /* 1. Initialise the model builder with the number of variables
            and the number of degrees of freedom in the DAE system. */
454
455
    csModelBuilder_t mb;
     uint32_t Nvariables = ...;
456
    uint32<sup>t</sup> Ndofs
                            = ...
457
    mb.Initialize_DAE_System(Nvariables, Ndofs);
458
459
     /* 2. Create and set model equations using the provided objects. */
460
                                        time = mb.GetTime();
x = mb.GetVariables();
     const csNumber_t&
461
     const std::vector<csNumber_t>& x
462
     const std::vector<csNumber_t>& dx_dt = mb.GetTimeDerivatives();
463
     const std::vector<csNumber_t>& y
464
                                                = mb.GetDegreesOfFreedom();
465
466
     std::vector<csNumber_t> equations(Nvariables);
     for(uint32_t i = 0; i < Nvariables; i++)
equations[i] = F(x, dx_dt, y, time);</pre>
467
                                                         <--- MODEL SPECIFIC CODE
468
469
    mb.SetModelEquations(equations);
470
471
     /* 3. Set the initial conditions. */
472
    std::vector<real_t> x0(Nvariables, 0.0);
473
     for(uint32_t i = 0; i < Nvariables; i++)</pre>
474
          x0[i] = ...;
                                                        <--- MODEL SPECIFIC CODE
475
476
     mb.SetVariableValues(x0);
477
478
     /* 4. Generate Compute Stack models by partitioning the DAE system.
479
480
           In this case a single model for a sequential simulation is generated. */
    /* 4.1 Specify the output directory and simulation options. */
std::string inputFilesDirectory = "...";
std::string simulationOptions = "...";
481
482
483
484
    /* 4.2 Instantiate the graph partitioner. */
csGraphPartitioner_Simple partitioner;
485
486
487
     /* 4.3 Partition the DAE system to generate a single Compute Stack model. */
488
    std::vector<csModelPtr> cs_models = mb.PartitionSystem(1, &partitioner);
489
490
     /* 5. Export the model(s) into a specified directory (or use them directly). */
491
    mb.ExportModels(cs_models, inputFilesDirectory, simulationOptions);
493
```

**Listing 2.** Graph partitioning for simulation on distributed memory systems

```
494
495
     /* 4.2 Instantiate METIS graph partitioner. */
     csGraphPartitioner_Metis partitioner(PartGraphRecursive);
496
497
     /* Change the input arguments of the partitioning algorithm. */
498
     /* 4.2.1 Specify the load balancing constraints (optional). */
std::vector<std::string> balancingConstraints = {"Ncs", "Nflops"};
499
500
501
502
     /* 4.2.2 Set the METIS partitioner options (optional). */
     std::vector<int32_t> options = partitioner.GetOptions(); /* default values */
503
     options[METIS_OPTION_NITER]
                                          = 10;
504
     options[METIS_OPTION_UFACTOR] = 30;
505
     partitioner.SetOptions(options);
506
507
     /* 4.2.3 Specify the number of FLOPs for mathematical operations (optional). */
508
     std::map<csUnaryFunctions,uint32_t> unaryOperationsFlops;
unaryOperationsFlops[eSqrt] = 12; /* i.e. the sqrt function requires 12 FLOPs */
binaryOperationsFlops[eDivide] = 6; /* i.e. the operator / requires 6 FLOPs */
509
510
511
512
513
     /* 4.3 Partition the system to generate Npe models (one per processing element). */
514
     std::vector<csModelPtr> cs_models = mb.PartitionSystem(Npe, &partitioner,
515
                                                                           balancingConstraints,
516
                                                                           true,
517
                                                                           unaryOperationsFlops,
518
                                                                           binaryOperationsFlops);
528
```

#### 521 Model exchange and parallel evaluation of model equations

The main goal of the OpenCS framework is specification of large scale equation-based models 522 for simulation on shared and distributed memory systems. In addition, the developed models can 523 also be used for model-exchange and for parallel evaluation of model equations (to improve the 524 simulation performance in existing simulators). The Compute Stack Differential Equations Model 525 is used for loading a model into a host simulator and as a common interface to the data required 526 for integration in time by ODE/DAE solvers (i.e. evaluation of equations and derivatives). The 527 procedure is identical in both cases and includes the following steps (source code listing 3): 528 Step 1. Initialise MPI. 529 **Step 2.** Instantiate the *csDifferentialEquationModel* object (a reference implementation of the 530 csDifferentialEquationModel\_t interface). 531 **Step 3.** Load the model from the specified directory with input files (or use the existing Compute 532 Stack Models directly). 533 **Step 4.** Instantiate and set the Compute Stack Evaluator. In this example the OpenMP Compute 534 Stack Evaluator is used. It accepts the number of threads as an argument in its constructor. 535 If zero is specified, the default number of threads will be used (typically equal to the 536 number of cores). 537 **Step 5.** Obtain the necessary information from the model such as the number of variables, 538 variable names, types, absolute tolerances, initial conditions and the sparsity pattern in the 539 Compressed Row Storage (CRS) format. 540 **Step 6.** Evaluate model equations and derivatives (typically in a loop). 54 **Step 6.1** Set the current values of state variables and derivatives using the *SetAndSynchro*-542 *niseData* function. At this point, for simulations on message passing multiprocessors 543 the MPI interface will be used to exchange the adjacent unknowns between process-544 ing elements. 545 **Step 6.2** Evaluate equations residuals (for DAE problems) or a Right Hand Side (for ODE 546 problems) using the *EvaluateEquations* function. 547 **Step 6.3** Evaluate derivatives (the Jacobian matrix) using the *EvaluateJacobian* function. 548 Here, *csMatrixAccess* t is used as a generic interface to the sparse matrix storage in 549 linear solvers. *inverseTimeStep* is an inverse of the current step taken by the solver. 550 SetAndSynchroniseData should be called only before a call to the EvaluateEqua-551 tions function. It is assumed that a call to SetAndSynchroniseData has already been 552 performed and the current values set and exchanged between processing elements. 553 This is a typical procedure in ODE/DAE solvers where the model equations are 554 always evaluated first and then, if required, the derivatives evaluated and a precondi-555 tioner recomputed (in iterative methods) or the Jacobian matrix re-factored (in direct 556 methods). 557 Step 7. Free the resources allocated in the model and the evaluator. 558

559 **Step 8.** Finalise MPI.

Listing 3. Procedure for model exchange

```
560
    /* 1. Initialise MPI. */
561
    int rank;
562
    MPI_Init(&argc, &argv);
563
    MPI_Comm mpi_world = MPI_COMM_WORLD;
564
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
565
566
    /* 2. Instantiate the Compute Stack model. */
567
    csDifferentialEquationModel model;
568
569
    /* 3. Load the model from the specified directory with input files. */
570
571
    model.Load(rank, inputFilesDirectory);
572
    /* 4. Instantiate and set the Compute Stack Evaluator. */
573
    csComputeStackEvaluator_OpenMP evaluator(0);
574
    model.SetComputeStackEvaluator(&evaluator);
575
576
    /* 5. Get the model information (i.e. the sparsity pattern in CRS format). */
577
                     N, Nnz;
578
    int
579
    std::vector<int> IA, JA;
    model.GetSparsityPattern(N, Nnz, IA, JA);
580
581
    /* 6. Evaluate model equations and derivatives (typically in a loop). */
582
          6.1 Set the current values of state variables and derivatives. */
583
    model.SetAndSynchroniseData(time, x, dx_dt);
584
585
           6.2 Evaluate residuals/Right Hand Side. */
586
    /*
    model.EvaluateEquations(time, residuals);
587
588
          6.2 Evaluate derivatives (the Jacobian matrix). */
589
    1+
    model.EvaluateJacobian(time, inverseTimeStep, ma);
590
591
    /* 7. Free the resources allocated in the model and the evaluator. */
592
593
    model.Free();
594
     /* 8. Finalise MPI. */
595
    MPI Finalize();
599
```

#### 598 Simulation on shared memory systems

Simulation on shared memory systems is performed by embedding a simulation into a host simulator or using a standalone OpenCS simulator (*csSimulator*). Simulations using the standalone *csSimulator* are performed by executing the simulator with a single argument specifying the directory with input files. Embedded simulations are started using the OpenCS *cs::Simulate* function (the source code listing 4) or, if a user-defined schedule is required, using the OpenCS simulation API (the source code listing 5). In the latter case, the procedure includes the following steps:

- 606 **Step 1.** Initialise MPI.
- <sup>607</sup> Step 2. Load the simulation\_options.json and get run-time options.
- <sup>608</sup> **Step 3.** Instantiate model, simulation and ODE/DAE solver objects.
- 609 Step 4. Load the model from the input directory.
- Step 5. Create and set the Compute Stack Evaluator. The application-specific evaluator can
   be instantiated or the information about the type of evaluator and its parameters can be
   obtained from the "Model.ComputeStackEvaluator" section.
- 613 **Step 6.** Initialise the simulation.
- <sup>614</sup> **Step 7.** Calculate corrected initial conditions at time = 0 (for DAE systems only).
- Step 8. Run the simulation using the default Run function or implement a custom schedule using
   the functions *Integrate*, *IntegrateForTimeInterval* and *IntegrateUntilTime* provided by the
   *daeSimulation\_t* class.
- 618 Step 9. Print the solver stats and finalise the simulation.

- <sup>619</sup> **Step 10.** Free the resources allocated in the model and the evaluator.
- 620 **Step 11.** Finalise MPI.

621

637

**Listing 4.** Simulation on shared memory systems using the *Simulate* function

```
/* 1. Initialise MPI. */
622
623
    MPI_Init(&argc, &argv);
624
    /* 2a. Run simulation using the input files from the specified directory: */
625
    cs::Simulate(inputFilesDirectory);
626
627
    /* 2b. Run simulation using an existing model: */
628
    csModelPtr model;
                                          // model is generated by the Model Builder
629
    std::string simulationOptions = ...; // options in JSON format
630
    std::string outputDirectory = ...; // a directory to store the simulation outputs
631
    cs::Simulate(model, simulationOptions, outputDirectory);
632
633
    /* 3. Finalise MPI. */
634
    MPI_Finalize();
635
```

**Listing 5.** Simulation on shared memory systems using the simulation API

```
/* 1. Initialise MPI. */
638
    MPI_Init(&argc, &argv);
639
640
    int rank;
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
641
642
    std::string inputFilesDirectory = "...";
643
644
    /* 2. Load the simulation_options.json and get run-time options. */
645
    std::string simulationOptionsFile = inputFilesDirectory + "/simulation_options.json";
646
    daeSimulationOptions& cfg = daeSimulationOptions::GetConfig();
647
648
    cfg.Load(simulationOptionsFile);
649
    std::string outputDirectory = cfg.GetString("Simulation.OutputDirectory");
650
651
    /* 3. Instantiate model, simulation and ODE/DAE solver objects. */
652
    daeModel t
653
                     model;
    daeSolver t
                     daesolver:
654
655
    daeSimulation_t simulation;
656
    /* 4. Load the model from the input directory. */
657
658
    model.Load(rank, inputFilesDirectory);
659
    /* 5. Create and set the Compute Stack Evaluator
660
    * (in general, using the data from the "Model.ComputeStackEvaluator" section). */
csComputeStackEvaluator_Sequential evaluator;
661
662
    model.SetComputeStackEvaluator(&evaluator);
663
664
665
    /* 6. Initialise the simulation. */
    simulation.Initialize(&model,
666
                            &daesolver,
667
                            startTime, timeHorizon, reportingInterval,
668
669
                            outputDirectory);
670
    /* 7. Calculate corrected initial conditions at time = 0 (DAE systems only). */
671
672
    simulation.SolveInitial();
673
    /* 8. Run the simulation using the default Run function
674
675
     *
          (or implement a custom schedule). */
    simulation.Run();
676
677
    /* 9. Print the solver stats and finalise the simulation. */
678
    simulation.PrintStats();
679
    simulation.Finalize();
680
681
    /* 10. Free the resources allocated in the model and the evaluator. */
682
    model.Free();
683
684
    /* 11. Finalise MPI. */
685
    MPI_Finalize();
689
```

#### Simulation on distributed memory systems

Simulation on distributed memory systems is typically performed using the standalone *csSimulator* and the parallel jobs are started using the commands specific to a particular implementation of the MPI standard. Examples for three different operating systems (GNU/Linux, Windows and macOS) and MPI implementations are given in the source code listing 6. Details on the available options for starting the parallel jobs can be found in the documentation of a particular implementation.

Listing 6. Simulation on distributed memory systems using the standalone OpenCS simulator

```
695
696
    # 1. GNU/Linux (i.e. using OpenMPI)
697
    # On a local machine:
    $ mpirun -np <Npe> csSimulator "inputFilesDirectory"
698
    # On multiple nodes:
699
    $ mpirun --hostfile <hostfilename> -np <Npe> csSimulator "inputFilesDirectory"
700
701
    # 2. Windows (i.e. using MS-MPI):
702
    # On a local machine:
703
    $ mpiexec /np <Npe> csSimulator "inputFilesDirectory"
704
    # On multiple nodes:
705
    $ mpiexec /gmachinefile <hostfilename> /np <Npe> csSimulator "inputFilesDirectory"
706
707
    # 3. macOS (i.e. using MPICH):
708
    # On a local machine:
709
    $ mpiexec -n <Npe> csSimulator "inputFilesDirectory"
710
    # On multiple nodes:
711
    $ mpiexec -f <hostfilename> -n <Npe> csSimulator "inputFilesDirectory"
713
```

#### 714 **APPLICATIONS**

#### <sup>715</sup> Case 1: transient two-dimensional diffusion-reaction equations, uniform grid

The model describes the process of auto-catalytic chemical reaction with oscillations known as the Brusselator PDE. The net reaction is  $A + B \rightarrow D + E$  with transient appearance of intermediates X and Y, where A and B are reactants and D and E are products (Strogatz, 1994). The model is originally implemented using SUNDIALS IDAS suite (Serban and Hindmarsh, 2016). Under conditions where components A and B are in vast excess during the chemical reaction the system dynamics is described by the following equations:

$$\frac{du}{dt} = k_1 \left( \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right) + R_u(u, v, t)$$

$$\frac{dv}{dt} = k_2 \left( \frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right) + R_v(u, v, t)$$
(1)

where the reaction rates  $R_u$  and  $R_v$  are defined as:

$$R_{u}(u,v,t) = u^{2}v - (B+1)u + A$$

$$R_{v}(u,v,t) = -u^{2}v + Bu$$
(2)

Here,  $k_1$  and  $k_2$  are diffusion constants, *A* and *B* are the concentrations of components A and B, and *u* and *v* are concentration of intermediaries X and Y. The equations are distributed on the square domain  $x \in [0, 10]$  and  $y \in [0, 10]$  and discretised by central differencing on a uniform 1000x1000 spatial mesh resulting in 2,000,000 unknowns. The boundary conditions are homogeneous Neumann (no normal flux at boundaries). The initial conditions are given by:  $u(x, y, t_0) = 1.0 - 0.5 \cos(\pi y)$  and  $v(x, y, t_0) = 3.5 - 2.5 \cos(\pi x)$ . The concentrations of <sup>729</sup> components A and B and the diffusion constants are held constant ( $A = 1, B = 3.4, k_1 = k_2 =$ <sup>730</sup> 0.002). The relative and absolute tolerances for all unknowns are set to 10<sup>-5</sup>. The system is <sup>731</sup> simulated for 10 seconds and the outputs are taken every 0.1 second. The C++ source code and <sup>732</sup> the compilation and usage instructions are given in the Supplemental Listing S10 and on the <sup>733</sup> OpenCS website (dae\_example\_3, http://www.daetools.com/opencs-tutorials.html). <sup>734</sup> Five different runs have been performed (Table 2). The first four runs are simulated on a <sup>735</sup> single CPU Model equations are evaluated sequentially in run C1-SEO using the OpenMP API

single CPU. Model equations are evaluated sequentially in run C1-SEQ, using the OpenMP API 735 in C1-OMP (8 core CPU), and using the OpenCL framework in C1-CL (single device setup, all 736 equations evaluated on a GPU) and C1-CLx2 (heterogeneous CPU/GPU setup where 70% of 737 equations are evaluated on GPU and 30% on CPU). In C1-MPI run the system is partitioned 738 by dividing the 2D mesh into eight quadrants, the simulation carried out on 8 CPUs using MPI 739 interface and equations evaluated sequentially in every processing element. The DAE system 740 is integrated in time using the variable-step variable-order backward differentiation formula 741 from SUNDIALS IDAS solver (Hindmarsh et al., 2005). Systems of linear equations are solved 742 using the SUNDIALS generalised minimal residual solver (GMRES) and the IFPACK (Sala 743 and Heroux, 2005) ILU preconditioner from Trilinos suite (in the original IDAS model the 744 band-block-diagonal preconditioner has been applied). The input parameters for the IFPACK 745 preconditioner are given in Table 3 where k is the fill-in factor,  $\alpha$  is the absolute threshold,  $\rho$  is 746 the relative threshold and  $\omega$  is the relax value. The simulations are carried out in 64-bit Debian 747 Stretch GNU/Linux and compiled using the gcc 6.3 compiler, OpenCS 1.1.0, MPI-3.1 from the 748 Open MPI v2.0.2 package, OpenMP 4.5 from the GOMP library, and OpenCL 1.2 from NVidia 749 CUDA 9.0 with v384.90 display driver. The hardware configuration consists of Intel i7-6700HQ 750 CPU (4 cores/8 threads at 2.6 GHz, 8 GB of RAM, 34.32 GFLOPs peak double precision) and 751 a discrete NVidia GeForce GTX 950M GPU (640 execution units at 914 MHz, 2 GB of RAM, 752 36.56 GFLOPs peak double precision). 753

 Table 2. Case 1: Simulation runs

Run	Simulation	Evaluation of model equations
C1-SEQ	1 CPU	Sequential
C1-OMP	1 CPU	OpenMP (8 threads)
C1-CL	1 CPU	Single device OpenCL (100% on GPU)
C1-CLx2	1 CPU	Heterogeneous OpenCL (70% on GPU, 30% on CPU)
C1-MPI	8 CPUs	Sequential on every PE

**Table 3.** Case 1: IFPACK ILU preconditioner parameters

Run	k	ρ	α	ω
C1-SEQ	3	1.0	0.1	0.5
C1-OMP	3	1.0	0.1	0.5
C1-CL	3	1.0	0.1	0.5
C1-CLx2	3	1.0	0.1	0.5
C1-MPI	1	1.0	0.1	0.0

# Case 2: transient two-dimensional convection-diffusion-reaction equations, uni form grid

The model describes the Chapman mechanism for ozone kinetics arising in athmospheric simulations (Schiesser and Lapidus, 1976). The reaction involves three components: ozone singlet (O), ozone ( $O_3$ ) and oxygen ( $O_2$ ), where the first two reactions are photo-chemical and contain diurnal rate coefficients. The model is originally presented in Wittman (1996) and implemented using SUNDIALS CVodes suite (Serban and Hindmarsh, 2015). The system dynamics is described by the following equations:

$$\frac{dc_i}{dt} = K_h \frac{\partial^2 c_i}{\partial x^2} + \frac{\partial}{\partial y} \left( K_\nu(y) \frac{\partial c_i}{\partial y} \right) + V \frac{\partial c_i}{\partial x} + R_i(c_1, c_2, t), \ i = 1, 2$$
(3)

where the reaction rates  $R_1$  and  $R_2$  are given by:

$$R_1(c_1, c_2, t) = -q_1c_1c_3 - q_2c_1c_2 + 2q_3(t)c_3 + q_4(t)c_2$$

$$R_2(c_1, c_2, t) = q_1c_1c_3 - q_2c_1c_2 - q_4(t)c_2$$
(4)

Here,  $c_1$ ,  $c_2$  and  $c_3$  are concentrations of O,  $O_3$  and  $O_2$ , respectively,  $q_1$ ,  $q_2$ ,  $q_3$  and  $q_4$  are reaction rate coefficients, V is velocity, and  $K_h$  and  $K_v$  are diffusion coefficients. The numerical values of the input parameters are:  $V = 10^{-3}$ ,  $K_h = 4 \cdot 10^{-6}$  and  $K_v(y) = 10^{-8} \exp(0.2y)$ .  $q_1$ ,  $q_2$  and  $c_3$  are constant ( $q_1 = 1.63 \cdot 10^{-16}$ ,  $q_2 = 4.66 \cdot 10^{-16}$ ,  $c_3 = 3.7 \cdot 10^{16}$ ) while  $q_3$  and  $q_4$  vary diurnally:

$$q_{3}(t) = \begin{cases} \exp\left(-A_{3}/\sin(\omega t)\right), & \text{if } \sin(\omega t) > 0\\ 0, & \text{otherwise} \end{cases}$$

$$q_{4}(t) = \begin{cases} \exp\left(-A_{4}/\sin(\omega t)\right), & \text{if } \sin(\omega t) > 0\\ 0, & \text{otherwise} \end{cases}$$
(5)

where  $\omega = \pi/43200$  and  $A_3$  and  $A_4$  are coefficients ( $A_3 = 22.62, A_4 = 7.601$ ). The equations are distributed on the square domain:  $x \in [0, 20]$  km and  $y \in [30, 50]$  km and discretised by central differencing on a uniform 1000x500 spatial mesh resulting in 1,000,000 unknowns. In the original CVodes model the equations were also discretised using the central differences, except for the advection term where a biased 3-point difference formula was used. The boundary conditions are homogeneous Neumann (no normal flux at boundaries). The initial conditions are given by:

$$c_{1}(x, y, t_{0}) = 10^{6} \alpha(x) \beta(y)$$

$$c_{2}(x, y, t_{0}) = 10^{12} \alpha(x) \beta(y)$$

$$\alpha(x) = 1 - (0.1(x - x_{mid}))^{2} + 0.5(0.1(x - x_{mid}))^{4}$$

$$\beta(y) = 1 - (0.1(y - y_{mid}))^{2} + 0.5(0.1(y - y_{mid}))^{4}$$
(6)

where  $x_{mid} = (0+20)/2 = 10$  and  $y_{mid} = (30+50)/2 = 40$  are mid points of the *x*, *y* domains. The relative and absolute tolerances for all unknowns are set to  $10^{-5}$ . The system is integrated for 86,400 seconds (1 day) and the outputs are taken every 100 seconds. The C++ source code and the compilation and usage instructions are given in the Supplemental Listing S11 and on the OpenCS website (ode\_example\_3, http://www.daetools.com/opencs-tutorials.html). Five different runs have been performed (Table 4) identical to those in Case 1. In the run C2-MPI the system is partitioned by dividing the 2D mesh into eight quadrants. The ODE system

<sup>780</sup> is integrated in time using the variable-step variable-order backward differentiation formula

available in SUNDIALS CVodes solver (Serban and Hindmarsh, 2005). Systems of linear
equations are solved using the SUNDIALS generalised minimal residual solver and the IFPACK
(Sala and Heroux, 2005) ILU preconditioner from Trilinos suite (in the original CVodes model
the 2x2 block-diagonal preconditioner has been applied). The input parameters for the IFPACK
preconditioner for all runs are given in Table 5. The simulations are carried out using the same
software and hardware as in Case 1.

Table 4.	Case 2:	Simulation	runs
----------	---------	------------	------

Run	Simulation	Evaluation of model equations
C2-SEQ	1 CPU	Sequential
C2-OMP	1 CPU	OpenMP (8 threads)
C2-CL	1 CPU	Single device OpenCL (100% on GPU)
C2-CLx2	1 CPU	Heterogeneous OpenCL (70% on GPU, 30% on CPU)
C2-MPI	8 CPUs	Sequential on every PE

**Table 5.** Case 2: IFPACK ILU preconditioner parameters

Run	k	ρ	α	ω
C2-SEQ	1	1.0	$10^{-5}$	0.0
C2-OMP	1	1.0	$10^{-5}$	0.0
C2-CL	1	1.0	$10^{-5}$	0.0
C2-CLx2	1	1.0	$10^{-5}$	0.0
C2-MPI	1	1.0	0.1	0.0

# 787 **RESULTS**

The numerical results are compared to the original SUNDIALS IDAS (Serban and Hindmarsh, 788 2016) and CVodes (Serban and Hindmarsh, 2015) models. Comparison of the concentration 789 u at the bottom-left point (x=0, y=0) between the OpenCS and the original IDAS model for 790 two different operating regimes are presented in Fig. 8 (stable regime for B = 1.7) and Fig. 9 791 (unstable regime for B = 3.4). Comparison of the concentration  $c_1$  at the bottom-left point (x=0, 792 y=30) between the OpenCS and the original CVodes model is presented in Fig. 10. 793 Four main and four sub-phases of the numerical solution have been analysed: 794 1. EvaluateEquations – evaluation of model equations (residuals or right-hand side) 795 2. LinearSystemSetup – setup of the linear equations solver, with two sub-phases: 796 2.1. EvaluateJacobian – evaluation of a Jacobian matrix 797 2.2. ComputePreconditioner – computation of a preconditioner using the Jacobian data 798 3. LinearSystemSolve – solution of a linear systems of equation, with two sub-phases: 799 3.1. ApplyPreconditioner – application of the preconditioner to solve the linear system 800 3.2. JacobianVectorProduct – Jacobian-vector multiplication, required in every iteration 801 of the linear solver (in SUNDIALS the difference quotient approximation is used 802 and requires an additional call to the EvaluateEquations function) 803 4. InterProcessDataExchange – exchange of adjacent unknowns between processing elements, 804 required before every call to EvaluateEquations. 805 The total integration time, the duration of individual phases of the numerical solution and the 806 percentage of the total integration time in individual phases are presented in Table 6 for Case 1 807 and Table 7 for Case 2. 808 The speed-ups of individual phases of the numerical solution, the maximum theoretical 809 overall speed-ups and the achieved overall simulation speed-ups are given in Table 8 for Case 1 810 and Table 9 for Case 2. The maximum theoretical speed-ups for evaluation of model equations 811 can be estimated using the maximum peak performance for individual platforms. For instance, 812 for C1-CL and C2-CL runs the theoretical speed-up is 36.56 GFLOPs/(34.32 GFLOPs/8 cores) 813 = 8.52, for C1-OMP and C2-OMP runs it is 8.00 (the number of cores), while for C1-CLx2 814 and C2-CLx2 runs it is 8.52 (GPU) + 8.00 (CPU) = 16.52. The maximum theoretical overall 815 simulation speed-ups can be calculated from the Amdahl's law using the data from Table 6 and 816 7 and the maximum peak performance for individual platforms: 1/(1-p+p/s), where p is 817 the portion of the solution that can be parallelised and s is the maximum theoretical speed-up 818 for evaluation of model equations. For runs that utilise OpenMP and OpenCL (only the model 819 equations and derivatives are evaluated in parallel) they are: (a) 2.02 for C1-OMP and 3.78 820 for C2-OMP, (b) 2.04 for C1-CL and 3.87 for C2-CL, and (c) 2.19 for C1-CLx2 and 4.75 for 821 C2-CLx2. The maximum theoretical overall speed-up for the MPI runs can be estimated assuming 822 that evaluation of model equations and the linear solution can be parallelised: 4.61 for Case 1 823 and 6.14 for Case 2. 824



**Figure 8.** Case 1: Plot of the concentration *u* at the bottom-left point (x=0, y=0) - stable regime (B = 1.7)



**Figure 9.** Case 1: Plot of the concentration *u* at the bottom-left point (x=0, y=0) - unstable regime (B = 3.4)



**Figure 10.** Case 2: Plot of the concentration  $c_1$  at the bottom-left point (x=0, y=30)

	C1-S	EQ	C1-0	MP	C1-0	CL	C1-C	Lx2	C1-N	/IPI
Phase	Time, s	%								
EvaluateEquations	211.96	20.53	60.22	11.43	69.56	13.26	54.37	11.27	61.64	15.06
LinearSystemSetup (total)	65.14	6.31	39.06	7.41	33.31	6.35	33.83	7.02	9.71	2.37
EvaluateJacobian	29.85	2.89	9.88	1.87	6.45	1.23	7.16	1.49	6.47	1.58
ComputePreconditioner	35.29	3.42	29.18	5.54	26.86	5.12	26.67	5.53	3.24	0.79
LinearSystemSolve (total)	646.71	62.63	340.19	64.54	339.32	64.69	311.72	64.64	248.43	60.69
ApplyPreconditioner	192.46	18.63	158.78	30.12	146.63	27.96	145.70	30.22	56.88	13.89
JacobianVectorProduct	355.12	34.39	100.91	19.07	116.54	22.21	91.09	19.07	108.38	26.48
InterProcessDataExchange	-	-	-	-	-	-	-	-	15.54	3.81
Integration (total)	1023.87		522.06		520.05		477.76		407.44	

**Table 6.** Case 1: Execution times for individual phases of the numerical solution

**Table 7.** Case 2: Execution times for individual phases of the numerical solution

	C2-S	EQ	C2-0	MP	C2-0	CL	C2-C	Lx2	C2-N	/IPI
Phase	Time, s	%								
EvaluateEquations	570.70	29.57	132.84	20.39	132.58	21.73	95.58	17.58	85.39	17.51
LinearSystemSetup (total)	319.24	16.54	114.81	17.62	76.43	12.53	96.98	17.84	66.47	13.63
EvaluateJacobian	270.21	14.00	72.31	11.10	33.69	5.52	52.52	9.66	52.61	10.79
ComputePreconditioner	49.03	2.54	42.50	6.52	42.74	7.01	44.46	8.18	13.86	2.84
LinearSystemSolve (total)	956.58	49.56	328.86	50.48	325.11	53.29	275.08	50.59	287.82	59.03
ApplyPreconditioner	83.15	4.31	73.67	11.30	70.30	11.52	72.66	13.37	40.41	8.29
JacobianVectorProduct	781.10	40.47	181.82	27.91	181.78	29.79	131.08	24.11	173.82	35.66
InterProcessDataExchange	-	-	-	-	-	-	-	-	11.70	2.39
Integration (total)	1930.08		651.47		610.11		543.66		487.54	

**Table 8.** Case 1: Speed-ups for individual phases of the numerical solution

Phase	C1-OMP	C1-CL	C1-CLx2	C1-MPI
EvaluateEquations	3.52	3.05	3.90	3.36
EvaluateJacobian	3.02	4.63	4.17	4.36
ComputePreconditioner	-	-	-	10.29
ApplyPreconditioner	-	-	-	3.40
JacobianVectorProduct	3.52	3.05	3.90	3.36
Max. theoretical overall speed-up	2.02	2.04	2.19	4.61
Overall speed-up	1.96 (97%)	1.97 (96%)	2.17 (98%)	2.51 (55%)

Phase	C2-OMP	C2-CL	C2-CLx2	C2-MPI
EvaluateEquations	4.30	4.30	5.97	4.18
EvaluateJacobian	3.74	8.02	5.14	3.72
ComputePreconditioner	-	-	-	2.56
ApplyPreconditioner	-	-	-	1.66
JacobianVectorProduct	4.30	4.30	5.96	4.18
Max. theoretical overall speed-up	3.78	3.87	4.75	6.14
Overall speed-up	2.96 (78%)	3.16 (82%)	3.55 (75%)	3.96 (64%)

Table 9. Case 2: Speed-ups for individual phases of the numerical solution

# **DISCUSSION**

<sup>826</sup> Comparison of the numerical results between the OpenCS model and the original SUNDIALS
<sup>827</sup> IDAS (for Case 1, Fig. 8 and 9) and CVodes (for Case 2, Fig. 10) models shows a good
<sup>828</sup> agreement. The observed small variations can be attributed to the internal implementation details:
<sup>829</sup> SUNDIALS models use different (and less efficient) preconditioners and, in addition, a different
<sup>830</sup> discretisation method has been applied to the advection term in Case 2.

The overall performance is in the following order for both cases: sequential < OpenMP <831 OpenCL (GPU) < OpenCL (CPU+GPU) < MPI implementations (Table 8 and 9) and agree well 832 with the theoretical limits. The achieved overall simulation speed-ups in Case 1 are 1.96, 1.97, 833 2.17 and 2.51 for C1-OMP, C1-CL, C1-CLx2 and C1-MPI runs, respectively (97, 96, 98 and 834 55% of the maximum theoretical overall speed-up, respectively). In Case 2 the achieved overall 835 simulation speed-ups are 2.96, 3.16, 3.55 and 3.96 for C2-OMP, C2-CL, C2-CLx2 and C2-MPI 836 runs, respectively (78, 82, 75 and 64% of the maximum theoretical overall speed-up, respectively). 837 In the OpenMP and the OpenCL runs the simulation is carried out on a single processor and 838 only evaluation of model equations is parallelised. Here, the OpenCL implementation performs 839 faster since the NVidia GPU device offers a higher maximum peak performance (36.56 GFLOPs) 840 than the eight Intel cores (34.32 GFLOPs). The reason for somewhat low overall speed-ups in 841 single processor simulations (especially in Case 1) is that both cases are dominated by the time 842 for solution of linear systems (62.63% of the total integration time in C1-SEQ and 49.56% in 843 C2-SEQ run, Table 6 and 7). The main reason is a costly Jacobian-vector multiplication phase 844 required in every iteration of the linear solver: the SUNDIALS GMRES solver uses a difference 845 quotient approximation of the Jacobian and requires 34.39% of the total integration time in 846 C1-SEQ and 40.47% in C2-SEQ run. As expected, the best performance is achieved in C1-MPI 847 and C2-MPI runs where the whole system is partitioned into eight sub-systems and independently 848 simulated. Again, the overall simulation speed-ups are lower than the maximum theoretical since 849 not all phases of the numerical solution can be parallelised, the performance of individual phases 850 of the numerical solution does not scale linearly with the number of processors and there is an 851 additional cost for inter-process data exchange between the processing elements (during the linear 852 algebra operations and before every evaluation of model equations). 853 The speed-ups in the EvaluateEquations phase are 3.52, 3.05, 3.90 and 3.36 for C1-OMP, 854

The speed-ups in the EvaluateEquations phase are 3.52, 3.05, 3.90 and 3.36 for C1-OMP, C1-CL, C1-CLx2 and C1-MPI runs, respectively (Table 8) and 4.30, 4.30, 5.97 and 4.18 for C2-OMP, C2-CL, C2-CLx2 and C2-MPI runs, respectively (Table 9). The speed-ups in the EvaluateJacobian phase are 3.02, 4.63, 4.17 and 4.36 for C1-OMP, C1-CL, C1-CLx2 and C1-MPI runs, respectively (Table 8) and 3.74, 8.02, 5.14 and 3.72 for C2-OMP, C2-CL, C2-CLx2 and

C2-MPI runs, respectively (Table 9). The maximum theoretical speed-up is 8.00 for OpenMP 859 and MPI runs, 8.52 for OpenCL runs and 16.52 for heterogeneous OpenCL. In Nikolić (2018) 860 it has been shown that the achieved speed-ups are in general higher in these two phases. The 861 reason is that much more complex expressions for model equations arising from the finite element 862 discretisation are used than in the finite difference equations studied in this work. Thus, a much 863 larger amount of computation is required and the hardware is better utilised. The speed-ups for 864 evaluation of the Jacobian are higher than speed-ups for evaluation of equations since a larger 865 number of evaluations are required and the hardware is again better utilised. The speed-ups in 866 the ComputePreconditioner phase are 10.29 for C1-MPI run (Table 8) and 2.56 for C2-MPI run 867 (Table 9). The speed-ups in the ApplyPreconditioner phase are 3.40 for C1-MPI run (Table 8) 868 and 1.66 for C2-MPI run (Table 9). The time for inter-process data exchange is only 3.81% of the 869 total integration time in Case 1 and 2.39% in Case 2 and does not significantly affect the overall 870 performance. The fact is that partitioning of the overall ODE/DAE system into a number of 871 ODE/DAE sub-systems has the largest effect on the evaluation of model equations and derivatives 872 (eight times lower the number of equations in every PE) and to a lesser degree on the solution of 873 linear systems (although the linear systems are eight times smaller in every PE the performance 874 does not scale linearly with the size of the problem). 875

# 876 CONCLUSIONS

The main ideas, the key concepts, the components, the algorithms and the API of the OpenCS framework are presented in this work. OpenCS provides an universal platform for modelling of problems described by systems of differential and algebraic equations, parallel evaluation of model equations on diverse types of computing devices, model exchange and parallel simulation on shared and distributed memory systems (including heterogeneous systems).

The framework offers the numerous benefits. For instance, model equations are transformed 882 into the postfix notation expression stacks, stored as an array of binary data and evaluated using 883 a stack machine. This way, the equations can be evaluated on virtually all computing devices 884 with no additional processing (including heterogeneous systems) and switching to a different 885 computing device is controlled by an input parameter. The OpenMP API is used for parallel 886 evaluation on general purpose processors and the OpenCL framework is used for parallelisation 887 on streaming processors and heterogeneous systems. The low-level model specification data 888 structures, stored as files in binary format, are used as an input for parallel simulations on all 889 platforms and provide a simple platform-independent binary interface for model exchange. The 890 partitioning algorithm can accurately balance the computation and memory loads in all important 891 phases of the numerical solution. Since a single simulation software and a common model-892 specification are utilised on all platforms, OpenCS models can be used for benchmarks between 893 different simulators, ODE/DAE solvers, individual computing devices and high performance 894 computing systems. 895

The capabilities of the framework are illustrated using two large scale problems. The overall performance and the performance of four main and four sub-phases of the numerical solution have been analysed. For simulations carried out on a single processor the OpenMP API and the OpenCL frameworks have been utilised for parallelisation of model equations. The MPI interface has been used for simulation on message-passing multiprocessors. It has been observed that the overall simulation performance is in the following order: sequential < OpenMP < OpenCL (single device) < OpenCL (heterogeneous CPU+GPU) < MPI implementations. As it has been expected, the MPI simulations offer the best performance since the whole ODE/DAE system is partitioned into a specified number of ODE/DAE sub-systems and independently simulated. However, the overall simulation speed-up is lower than the maximum theoretical: the performance of individual phases of the numerical solution does not scale linearly with the number of processors and there is an additional cost for inter-process data exchange between the processing elements.

The future work will focus on applications of the framework to large-scale multi-scale and multi-physics problems, further improvement of the performance and reduction of the memory requirements, implementation of problem-specific graph partitioners, new solvers and preconditioner libraries, and Compute Stack Evaluator implementations for additional types of computing devices (such as Xeon Phi and FPGA).

# 913 **REFERENCES**

Johan Akesson, Karl-Erik Arzen, Magnus Gafvert, Tove Bergdahl, and Hubertus Tummescheit. Modeling and optimization with Optimica and JModelica.org - languages and tools for solving

large-scale dynamic optimization problems. *Comput. Chem. Eng.*, 34(11):1737–1749, 2010.

917 Altair. HyperWorks, 2018. URL https://altairhyperworks.com.

<sup>918</sup> Christian Andersson, Claus Fuhrer, and Johan Akesson. Assimulo: A unified framework for
<sup>919</sup> ode solvers. *Math. Comput. Simulat.*, 116(0):26 – 43, 2015. ISSN 0378-4754. doi: 10.1016/j.

matcom.2015.04.007. URL http://dx.doi.org/10.1016/j.matcom.2015.04.007.

921 Ansys, Inc. ANSYS Fluent, 2018. URL http://www.ansys.com.

<sup>922</sup> Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman,

Lisandro Dalcin, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley,

Lois Curfman McInnes, Karl Rupp, Barry F. Smith, Stefano Zampini, and Hong Zhang. PETSc

users manual. Technical Report ANL-95/11 - Revision 3.6, Argonne National Laboratory,

2015. URL http://www.mcs.anl.gov/petsc.

W. Bangerth, R. Hartmann, and G. Kanschat. deal.II – a general purpose object oriented finite
element library. *ACM Trans. Math. Softw.*, 33(4):24/1–24/27, 2007.

Paul I. Barton and Constantinos C. Pantelides. Modeling of combined discrete/continuous

processes. AIChE J., 40(6):966–979, 1994. ISSN 1547-5905. doi: 10.1002/aic.690400608.
 URL http://dx.doi.org/10.1002/aic.690400608.

<sup>932</sup> COMSOL, Inc. COMSOL Multiphysics, 2018. URL http://www.comsol.com.

<sup>933</sup> Dassault Systemes. Abaqus, 2018. URL http://www.simulia.com.

Peter Fritzson and Vadim Engelson. Modelica — A unified object-oriented language for system

modeling and simulation. In Eric Jul, editor, *ECOOP'98 — Object-Oriented Programming*,

volume 1445 of *Lecture Notes in Computer Science*, pages 67–90. Springer Berlin, Heidelberg,

937 Germany, 1998. ISBN 978-3-540-64737-9. doi: 10.1007/BFb0054087. URL http://dx.doi.org/

938 10.1007/BFb0054087.

Peter Fritzson, Peter Aronsson, Hakan Lundvall, Kaj Nystrom, Adrian Pop, Levon Saldamli, and
 David Broman. The OpenModelica modeling, simulation, and development environment. 46th

<sup>941</sup> Conference on Simulation and Modelling on the Scandinavian Simulation Society (SIMS2005),

<sup>942</sup> Trondheim, Norway, 2005. URL http://scansims.org/sims2005/SIMS2005\_58.pdf.

John D. Hedengren, Reza Asgharzadeh Shishavan, Kody M. Powell, and Thomas F. Edgar.

Nonlinear modeling, estimation and predictive control in apmonitor. *Comput. Chem. Eng.*, 70:

<sup>945</sup> 133 – 148, 2014. ISSN 0098-1354. doi: 10.1016/j.compchemeng.2014.04.013. URL http:

- //www.sciencedirect.com/science/article/pii/S0098135414001306. Manfred Morari Special
   Issue.
- 948 Michael A. Heroux, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu,
- <sup>949</sup> Tamara G. Kolda, Richard B. Lehoucq, Kevin R. Long, Roger P. Pawlowski, Eric T. Phipps,
- Andrew G. Salinger, Heidi K. Thornquist, Ray S. Tuminaro, James M. Willenbring, Alan
- Williams, and Kendall S. Stanley. An overview of the trilinos project. ACM Trans. Math.
- 952 Softw., 31(3):397–423, 2005. ISSN 0098-3500. doi: 10.1145/1089014.1089021.
- Alan C. Hindmarsh, Peter N. Brown, Keith E. Grant, Steven L. Lee, Radu Serban, Dan E.
- <sup>954</sup> Shumaker, and Carol S. Woodward. SUNDIALS: Suite of Nonlinear and Differential/Algebraic
- <sup>955</sup> Equation Solvers. *ACM Trans. Math. Softw.*, 31(3):363–396, September 2005. ISSN 0098-3500.
- doi: 10.1145/1089014.1089020. URL http://doi.acm.org/10.1145/1089014.1089020.
- George Karypis and Vipin Kumar. A software package for partitioning unstructured graphs,
   partitioning meshes, and computing fill-reducing orderings of sparse matrices, 1995.
- 959 B. S. Kirk, J. W. Peterson, R. H. Stogner, and G. F. Carey. libMesh: A C++ Library for Parallel
- Adaptive Mesh Refinement/Coarsening Simulations. *Engineering with Computers*, 22(3–4): 237–254, 2006. https://doi.org/10.1007/s00366-006-0049-3.
- <sup>962</sup> Dragan D. Nikolić. DAE Tools: Equation-based object-oriented modelling, simulation and
- optimisation software. *PeerJ Computer Science*, 2:e54, April 2016. ISSN 2376-5992. doi:
- <sup>964</sup> 10.7717/peerj-cs.54. URL https://doi.org/10.7717/peerj-cs.54.
- Dragan D. Nikolić. Parallelisation of equation-based simulation programs on heterogeneous
   computing systems. *PeerJ Computer Science*, 4:e160, August 2018. ISSN 2376-5992. doi:
   10.7717/peerj-cs.160. URL https://doi.org/10.7717/peerj-cs.160.
- <sup>968</sup> Dragan D. Nikolić. Parallelisation of equation-based simulation programs on distributed memory
- systems. PeerJ Computer Science, 5:exxx, February 2019. ISSN 2376-5992. doi: 10.7717/
- peerj-cs.xxx. URL https://doi.org/10.7717/peerj-cs.xxx.
- P. C. Piela, T. G. Epperly, K. M. Westerberg, and A. W. Westerberg. ASCEND: an object-oriented computer environment for modeling and analysis: The modeling language. *Comput. Chem.*
- 973 Eng., 15(1):53-72, 1991. ISSN 0098-1354. doi: 10.1016/0098-1354(91)87006-U. URL
- http://www.sciencedirect.com/science/article/pii/009813549187006U.
- M. Sala and M. Heroux. Robust algebraic preconditioners with IFPACK 3.0. Technical Report
   SAND-0662, Sandia National Laboratories, 2005.
- W. E. Schiesser and L. Lapidus. Academic Press, Inc, New York, United States, 1976. ISBN
   0-12-4366406.
- 979 Radu Serban and Alan C. Hindmarsh. CVODES, the sensitivity-enabled ODE solver in SUN-
- DIALS. Proceedings of the 5th International Conference on Multibody Systems, Nonlinear
   Dynamics and Control, Long Beach, CA, USA, 2005.
- Radu Serban and Alan C. Hindmarsh. Example programs for CVODES v2.8.2. Technical Report
- UCRL-SM-208115, Center for Applied Scientific Computing, Lawrence Livermore National
   Laboratory, 2015. URL https://computation.llnl.gov/sites/default/files/public/cvs\_examples.
   pdf.
- <sup>986</sup> Radu Serban and Alan C. Hindmarsh. Example programs for IDAS v1.3.0. Technical Report
- <sup>987</sup> LLNL-TR-437091, Center for Applied Scientific Computing, Lawrence Livermore National
- Laboratory, 2016. URL https://computation.llnl.gov/sites/default/files/public/idas\_examples.
   pdf.
- Siemens. Multidisciplinary Design Exploration, 2018. URL https://mdx.plm.automation.siemens.
   com.

- <sup>992</sup> Steven H. Strogatz. CRC Press LLC, Boca Raton, Florida, USA, 1994. ISBN 0-73-8204536.
- <sup>993</sup> The MathWorks, Inc. MATLAB, 2018a. URL https://mathworks.com/products/matlab.
- <sup>994</sup> The MathWorks, Inc. Simulink, 2018b. URL https://mathworks.com/products/matlab.
- <sup>995</sup> The OpenFOAM Foundation. OpenFOAM, 2018. URL http://www.openfoam.org.
- <sup>996</sup> Waterloo Maple, Inc. Maple, 2015. URL http://www.maplesoft.com/products/maple/.
- <sup>997</sup> Michael R. Wittman. Testing of PVODE, a parallel ode solver. Technical Report UCRL-ID-
- 125562, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory,
   1996.
- 1000 Wolfram Research, Inc. Mathematica, 2015. URL https://www.wolfram.com/mathematica.